A Model for Generating Expectations: the Bridge Between Memory and Surprise

Luís Macedo
Departamento de Eng. Informática e de Sistemas – Instituto Superior de Engenharia de Coimbra / Centro de Informática e Sistemas da Universidade de Coimbra,
Quinta da Nora, Coimbra, Portugal
lmacedo@isec.pt

Amílcar Cardoso
Centro de Informática e Sistemas da Universidade de Coimbra, Departamento de Eng. Informática da Universidade de Coimbra,
Pinhal de Marrocos, Coimbra, Portugal
amilcar@dei.uc.pt

Abstract

In this paper we present a model for generating expectations in an artificial agent and evaluate through an experiment its connection both with the agent's memory (in this paper restricted to the episodic memory) and with surprise. Performed from two points of view, the creator and the evaluator, the experiment allows us to take conclusions about those connections and the role played by this model both in the production and evaluation of creative products. We discuss both the results of the experiment and the relation of our model of generating expectations with other models of creativity.

1 Introduction

Artificial Intelligence attempts to understand and build intelligent agents. Particularly, one of the approaches of Artificial Intelligence aims to understand and build artificial agents that act and think like humans [Russel and Norvig, 1995]. Roughly speaking, agents accept percepts from the environment and generate actions. They perform well when they do the "right" action. This process of making the right decision depends heavily on a good model of the environment that surrounds them. Unfortunately, the real world is not crystal clear to agents. Agents almost never have access to the whole environment, mainly because of the incompleteness and incorrectness of their perceptual and understanding components. Actually, it is too much work to obtain all the information from a complex and dynamic world, and it is quite likely that the accessible information suffers distortions. Nevertheless, since the success of agents depends heavily on the completeness of the information of the state of the world, they have to pursue alternatives to construct good models of the world even (and specially) when this is uncertain. According to psychologists, cognitive scientists and ethologists [Kline, 1999; Piaget, 1952], humans and, in general, animals attempt to overcome this limitation through the generation of *assumptions* or *expectations*[1] to fill in gaps in the present observational information. Note, however, that not all those expectations are made explicit. However, the reasoning of the agent may be improved if its model of the world contains also a good model of the future worlds. In this case, the process cannot be confined to fill in gaps in the information provided by perception because there is no information at all of those worlds. In order to overcome this limitation, agents also exhibit the ability to make predictions about future states of the world taking into account the present world and inference processes. Anyway, when the missing information, either of the present state of the world or of the future states of the world, becomes known to the agent, it might happen an inconsistence or conflict between it and the assumptions or expectations that the agent has. As defended by Reisenzein [2000], Gärdenfors [1994], Ortony and Partridge [1987], etc., the result of this inconsistence gives rise to surprise. It also gives rise to the process of updating beliefs, called belief revision (e.g., [Gärdenfors, 1992]).

Gärdenfors [1994] defended that expectations are defeasible beliefs that are necessary to everyday reasoning. With respect to their cognitive origins, Gärdenfors argued that they are much like summaries of previous experiences. Thus, he defended that they are the result of inductive reasoning.

[1] Although some authors use the terms assumption and expectation as synonyms, there are authors that make a distinction between them defending that an expectation has to do with future world states while assumptions are related with the current world state.

In their cognitive-psychoevolutionary model, the research group of the University of Bielefeld has considered some of these ideas, namely that surprise is elicited by the appraisal of unexpectedness. More precisely, it is proposed that surprise-eliciting events give rise to the following series of mental processes: (i) the appraisal of a cognized event as exceeding some threshold value of unexpectedness (schema-discrepancy) - according to Reisenzein [1999], this is achieved by a specialized comparator mechanism, the unexpectedness function, that computes the degree of discrepancy between "new" and "old" beliefs or schemas; (ii) interruption of ongoing information processing and reallocation of processing resources to the investigation of the unexpected event; (iii) analysis/evaluation of that event; (iv) possibly, immediate reactions to that event and/or updating or revision of the "old" schemas or beliefs.

Sharing some of these ideas, Ortony and Partridge [1987] proposed a model for surprise. The following assumptions were made in their model: the system (or agent) receives an input proposition; the system has an episodic and semantic memory; elements of the memory may be immutable (propositions that are believed to be always true) or typical (those that are believed to be sometimes true); and, some elements of the memory are activated when an input proposition is received. Moreover, they distinguished between *practically deducible propositions* and *practically non-deducible propositions*. *Practically deducible propositions* comprises the propositions that are explicitly represented in memory, as well as those that can be inferred from them by few and simple deductions. Hence, *practically deducible propositions* are that subset of formally deducible propositions that don't require many and complex inferences. Furthermore, *practically deducible propositions* may be actively or passively deduced in a particular context. In the former case, their content corresponds to *actively expected* or *predicted* events; in the latter case, to *passively expected* (*assumed*) events. Based on these assumptions, Ortony and Partridge proposed that surprise may result from three situations: (i) *active expectation failure*: here, surprise results from a conflict or inconsistency between the input proposition and an *active prediction* or *expectation*; (ii) *passive expectation failure* (or *assumption failure*): here, surprise results from a conflict or inconsistency between the input proposition and what the agent implicitly knows or believes (*passive expectations* or *assumptions*); and (iii) *unanticipated incongruities* or deviations from norms: here, surprise results from a conflict or inconsistency between the input proposition (which in this case is a *practically non-deducible proposition*) and what, after the fact, may be judged to be normal or usual [Kahneman and Miller, 1986], that is, *practically deducible propositions* (immutable or typical) that are suggested by the unexpected fact. Note that, in this case, at least prior to the unexpected event, there are no expectations (passive or active) with which the input proposition could conflict.

Within creativity, surprise and related concepts such as unpredictability or unexpectedness have been pointed as relevant characteristics of a creative product, in addition to other commonly referred features such as novelty, originality, interestingness and appropriateness (also defined as usefulness, aesthetic value, rightness, etc.) [Boden, 1992, 1995; Jackson and Messick, 1967; Koestler, 1964; Lubart, 1994; Macedo, 1998; MacKinnon, 1962; Moorman and Ram, 1994; Ritchie, 2001; Saunders and Gero, 2001]. Furthermore, Boden argued that there is a distinction between mere novelty and creativity [Boden, 1995]. In her opinion, that distinction resides on the fact that creative products are not only novel but also unpredictable, unexpected and therefore surprising. According to Boden, unpredictability is the essence of creativity: creative products amaze us, shock us and delight us mainly because they are unexpected or unpredictable. Besides these statements, according to Boden, a creative idea not only did not happen before, but mainly could not have happened before. Moreover, it has to be produced using generative rules different from those used to produce non-creative ideas. We showed to some extent that distinction between novelty and creativity stressing the relevance of unexpectedness/surprise in novel products so that they can be classified as creative [Macedo and Cardoso, 2002].

As a result of these ideas comes that, in order to build an artificial agent capable to evaluate or even produce creative products as humans do, there must be an adequate model for inferring practically deducible expectations (explicit knowledge plus inferred knowledge). This way, when the agent or system is faced with a particular state of the world it can generate those expectations as humans do in similar circumstances. This is as much important as the step of the process that follows expectation generation. This one involves the computation of the degree of conflict or inconsistence between the input knowledge (product, in the case of creativity) and the practically deducible knowledge.

In this paper, we focus on a model for the generation of expectations and its relation with the agent's memory and surprise.

The next section describes the model for generating expectations. Subsequently, we present an experimental test carried out to evaluate its connection with the agent's memory and surprise. Finally, we make a discussion and present some conclusions.

2 The Model for Generating Expectations

We developed a multi-agent environment in which, in addition to inanimate agents (objects such as buildings), there are two main kinds of animate, interacting agents: the *author-agents* or *creators*, whose main function is to create things (objects, events), and the *jury-agents* or *explorers* whose goal is to explore the environment by

analyzing, studying and evaluating it. An agent can also exhibit both characteristics.

The environment is previously defined by the user and comprises a variety of objects located at specific positions. In the present article, these objects are confined to buildings. Each object comprises three distinct, fundamental components: *structure*, *function* and *behavior* [Goel, 1992]. For the sake of simplicity, the structure (the visible part of the object), is restricted to the shape of the object (e.g., triangular, rectangular, etc.); however, any object may comprise several sub-objects. The function of the object concerns its role in the environment (e.g., house, church, hotel, etc.). The behavior of the object concerns its activity (actions and reactions) in response to particular features of external or internal stimuli (e.g., static, mobile) (for more details about the environment and characteristics of the agents see [Macedo and Cardoso, 2001c, 2002]).

The perceptual system of an agent (two simulated sensors) provides information related to the structure, the function, and the behavior of the objects, as well as the distance of the objects. The function of the objects is not accessible (i.e., cannot be inferred from visual information) unless an agent is at the same place as the object.

An agent's knowledge base is exclusively of an episodic kind (for an example, see Figure 1), rather than being both semantic and episodic in nature (although this will be part of our future work): each object is stored as an individual case in the episodic memory and associated with a number that expresses its absolute frequency (see Figure 1). In addition to this case-base, the agent also has a map of the environment in memory where it stores the location of the objects. This knowledge base comprising explicit representations of objects constitutes the *explicit knowledge*. Following Ortony and Partridge, we also distinguish between *deducible* and *non-deducible knowledge* (see Figure 2). Both categories constitute the *formally deduced knowledge*. Thus, using simple inference rules it is possible to produce an amount of knowledge, called *inferred knowledge*, that is not explicitly represented. Together with explicitly represented knowledge, it forms the category of *deducible knowledge*. *Non-deducible knowledge* corresponds to that knowledge that is not explicitly represented and that involves complex inference rules to deduce it, probably because it is not motivated or not easily explained from the current state of the world. As in Ortony and Partridge model, there is an *immutability* property associated to each element of knowledge (proposition). In our model, the immutability can be extracted from the absolute frequency values associated with the cases stored in episodic memory (see Figure 1). For instance, in the example shown in Figure 1, the proposition "houses have square facades" is immutable (since all the houses in memory have squared facades), whereas "houses have square windows" is a typical proposition with a probability (immutability) value of 0.50 (as implied by Ortony and Partridge's model, immutability is a continuous variable).

Field \ Case	C_1	C_2	C_3	C_4
Structure				
Function	House	House	Church	Hotel
Behavior	Static	Static	Static	Static
Abs. Freq.	50	40	5	5

Figure 1. Example of an episodic memory in the domain of buildings.

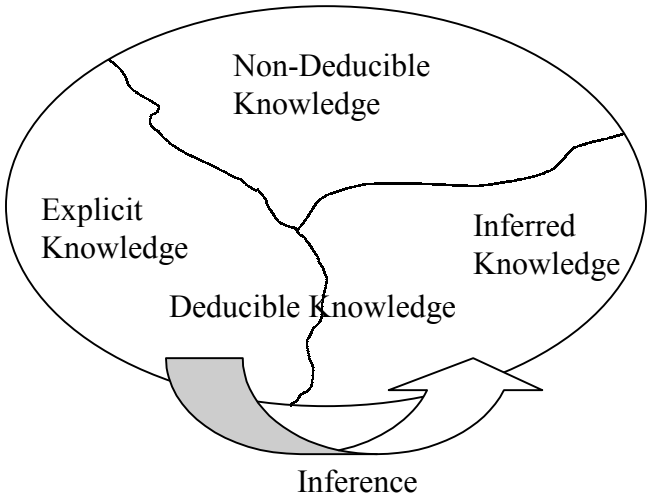

Figure 2. Categories of knowledge.

The usual activity of the agents consists in moving through the environment, hoping to find interesting things (objects or events) that deserve to be investigated. When the surrounding world (objects or events) is perceived, the agent "consciously" deduces knowledge for the missing information (e.g., "it is a house with 67% of probability", "it is a hotel with 45% of probability", etc.; note that the actual function of a building becomes available to the agent only when its position and that of the building are the same). In this case, the agent has performed active predictions or active expectations. In addition, it might already had knowledge or believes about the rest of the information (e.g., "doors are immutably squared", "windows are typically squared"). This means that the agent has performed passive assumptions or passive expectations. Both passive and active expectations are therefore within the category of deducible knowledge. We call *expectation set* to the set comprising passive and active expectations.

The new knowledge (input knowledge) provided by the new state of the world may be of two kinds from the perspective of the agent: (a) practically deducible knowledge, i.e., the new knowledge (the state of the world or portion of it) is easily inferred, explainable or motivated by the previous state of the world and from the agent's knowledge base, or (b) non-deducible knowledge, i.e.,

the new knowledge is hardly inferred, unexplainable or not motivated by the previous state of the world and from the agent's knowledge base. In the former case, (a), another two major situations may happen: (a1) there is a conflict between that new knowledge and the practically deduced knowledge of the agent (active and passive expectations), or (a2) there is no conflict at all. In the former case, (a1), the degree of the conflict is assessed and the respective reaction is performed such as the elicitation of surprise [Macedo and Cardoso, 2001b]. In situation (b), there is no conflict between the input knowledge and expectations because there are no expectations at all. Notice that the input knowledge of situation (b) describes a state of the world or part of it that is non-deducible to the agent and therefore it is impossible for the agent to have expectations for it. However, after the fact (after that unpredictable or practically non-deducible state of the world) it is possible to evaluate its deviation from the norm or the usual. Thus, in contrary to situation (a1), the conflict or inconsistence happens after the fact and is related to practically deducible knowledge that it suggests.

Although other probabilistic methods might be used to compute the probability of an expected knowledge element X, P(X), in case of objects comprising several components we may compute the probability of the whole object Obj_k computing the mean of the conditional probabilities of their n constituent parts, which are individually computed using the Bayes's formula [Shafer and Pearl, 1990] (notice that each one of those conditional probabilities individually gives the degree of expectedness of a specific piece of the object, given as evidence the rest of the object):

$$P(Obj_k)=\frac{\sum_{l=1}^{n} P(Obj_k^l \mid Obj_k^1, Obj_k^2, \ldots, Obj_k^{l-1}, Obj_k^{l+1}, \ldots, Obj_k^n)}{n}$$

Let us consider an example to illustrate this model for generating expectations. Consider the situation depicted in Figure 3 and assume that the episodic memory of the agent is the one presented in Figure 1. Thus, the deducible knowledge of the agent comprises not only this explicit representations of objects but also additional passive expectations deduced from it such as "every church has a cross", "hotels don't have doors nor windows", etc. After perceiving the surrounding environment (in this case three objects), the input knowledge is: "object 1 with a squared window, rectangular door, squared facade and triangular roof", "object 2 with circular window, rectangular door, squared facade and triangular roof", and "object 3 with a squared window, rectangular door, trapezoidal facade and triangular roof". Immediately, the agent computes active expectations for the missing information such as the function of each object (remember that the function of an object is accessible only when the position of the object and that of the agent are the same). For instance, the possible functions for object 1 are: house, church and hotel. Actually, these form the exhaustive list of the functions of objects stored in memory. Applying the conditional probability equation, it computes values for P(house), P(church) and P(hotel). Meanwhile, the passive expectations for the available information are ready to be compared with the available information. For instance, with respect to object 3, the agent would never expect a trapezoidal facade. Actually, there is no object in memory whose facade has such a shape. Therefore there is a conflict of this portion of knowledge with the deducible knowledge available in the agent's memory. Thus, there is a passive expectation failure. For the next step of the process, the computation of the degree of conflict and the computation of the intensity of surprise, we use of the following surprise function [Macedo and Cardoso, 2001b]:

$$SURPRISE(Agt, Obj_k) = DegreeOfUnexpectedness(Obj_k, Agt(Memory)) = 1 - P(Obj_k)$$

where P(Objk) is computed by the equation presented above.

With respect to object 1 and 2 and in the absence of their functions, there is no conflict at all. Actually, both buildings are entirely present in memory.

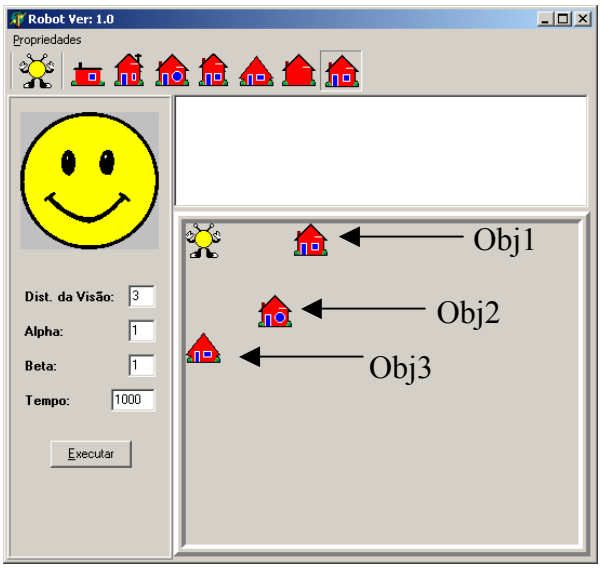

Figure 3. Example of a state of the world.

Consider now that the agent visits object 1. Once again the agent computes active expectations that together with its passive expectations and the active expectations computed in the previous state of the world form the deducible knowledge. Suppose the agent becomes aware that object 1 is a house. In this case there is no expectation violation since this function was the most expected one. In contrary, suppose the function of object 1 is church. In

this case there is an (active) expectation violation since this was not the most expected function.

3 Experimental Test

We performed an experiment to evaluate the connection between the episodic memory and the model for generating expectations and between this model and the surprise value of creative products. In order to achieve this goal, the following issues were tested: (i) whether the size of the episodic memory influences the size of the expectation set; (ii) whether the size of the expectation set influences both the evaluation and production of creative products in terms of surprise.

In the experiment, an author-agent was asked to produce the five most surprising products it could. This request was repeated eight times, each time with a larger episodic memory (see Figure 4), resulting in eight sets of buildings (denoted in Figures 4 by *A*, *B*, …, *H*), each one with five buildings. This means that the set of buildings *A* was produced using an episodic memory smaller than that used to produce the set of buildings *B*, and so on. The expectation set generated by the author-agent for their functions during this process of production was collected (see Figure 4).

Subsequently, those buildings were provided to a set of fifteen jury-agents, selected after exploring part of the environment. Since those agents had different ages, they had different knowledge levels, i.e., they had different episodic memories. These memories ranged from very small memories, with only 5 buildings, to large memories with 229 buildings, following an approximately linear distribution. For each set of objects produced by the author-agent, the average of the intensities of surprise elicited in the fifteen jury-agents was computed with respect to the function of the objects (see Figure 4). In addition, the expectation set of these jury-agents for the function of each object was collected.

The results of the experiment are presented from two points of view: the point of view of the author-agent (Figure 4) and the point of view of the jury-agents (Figure 5). The former point of view allows taking conclusions about the connection between the author-agent's episodic memory and its expectations and between these and the surprise-value of the products it creates. The latter point of view allows taking the same conclusions but with respect to the jury-agents, and therefore it relates expectations with the evaluation of creative products in terms of surprise.

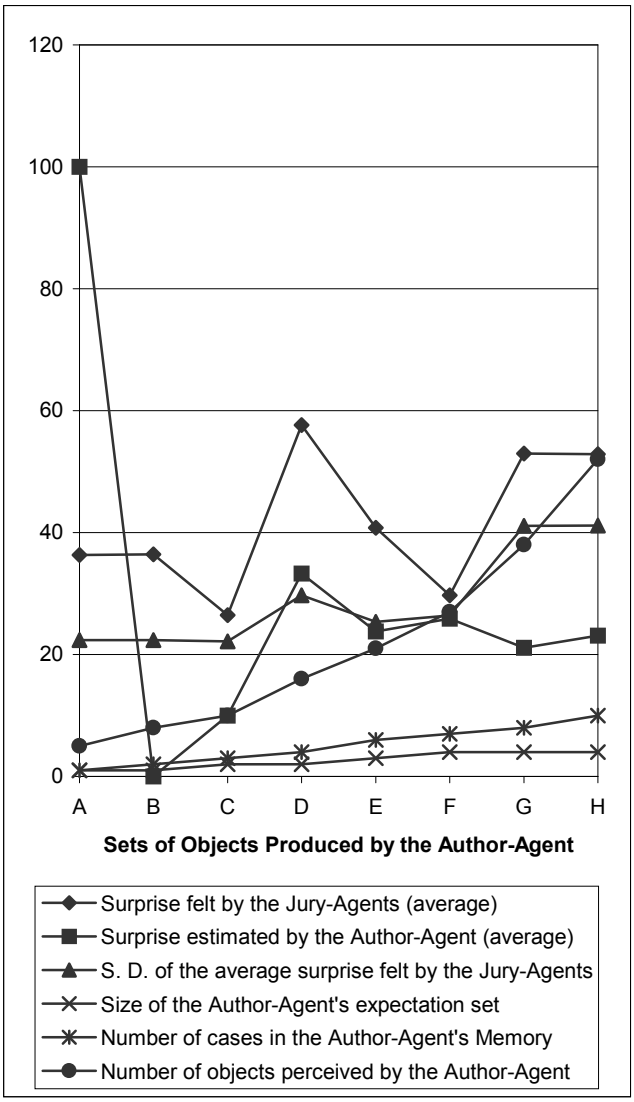

Figure 4. Results of the experiment from the perspective of the author-agent. Note that, unlike the other data, the intensity of surprise is measured in percentage.

According to the results of the experiment presented in Figure 4, we may conclude that the larger the episodic memory of an author-agent, the larger its expectation set is. However, as it can be seen in Figure 4, the size of this expectation set is more correlated with the number of cases in memory (correlation = 0.96) than with the number of buildings stored in memory (correlation = 0.89). These results also confirm the conclusions achieved before by ourselves [Macedo and Cardoso, 2001a]: on average, the larger the episodic memory of an author-agent, the higher the intensity of surprise elicited in the jury-agents by the products it creates. Since we consider surprise as a feature of creative solutions, we may infer that highly creative solutions are more likely achieved when the author-agent has a larger episodic memory. Therefore we may also conclude that the larger the expectation set,

the higher the intensity of surprise elicited in the jury-agents by the products it creates.

According to the results of the experiment presented in Figure 5, we may conclude also that the larger the episodic memory of a jury-agent, the larger its expectation set is. Once again, the size of this expectation set is more correlated with the number of cases in memory (correlation = 0.90) than with the number of buildings stored in memory (correlation = 0.63).

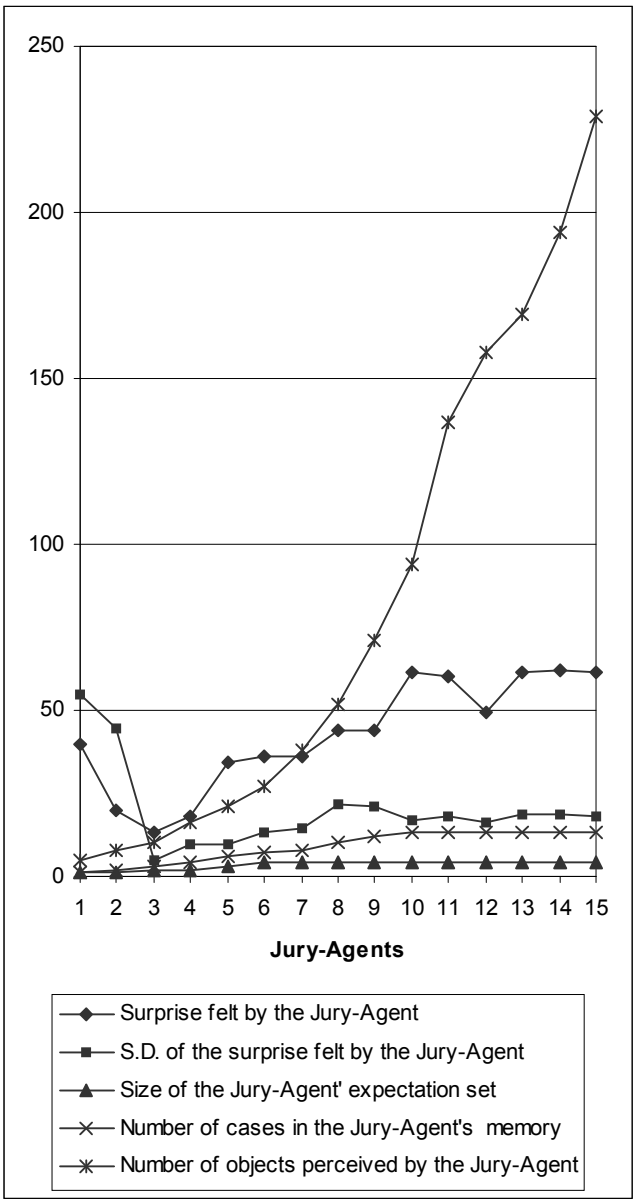

Figure 5. Results of the experiment from the perspective of the jury-agents. Again the intensity of surprise is measured in percentage.

4 Discussion

The results achieved in the experiment presented above may be explained as follows.

It is to some extent obvious that the larger the episodic memory, the higher the probability of the knowledge being being diverse. For instance, an agent with an episodic memory of 50 buildings probably has different kinds of objects such as houses, hotels, churches, etc., with a variety of structures. In contrary, an agent with two buildings can only have two kinds of structures associated with only at most two functions. Therefore, when faced with a new building, the agent with a higher diversity of knowledge probably expects several functions for it while the agent with a lower diverse episodic memory can only expect two functions at most. This may explain the high correlation between the size of the episodic memory and the size of the expectation set. However, the number of cases seems to be more important than the number of buildings perceived by an agent as confirmed by the correlation values presented above. Actually, when two or more similar buildings are perceived, they are represented by a single case associated with its frequency that statistically reflects that number of buildings it represents. Thus, the diversity of cases is more important to the generation of larger expectation sets than the amount of buildings perceived. This is deeply related with issues such as case-base competence (e.g., [Smyth and McKenna, 1999]). Notice that equal size does not mean equal competence because, for instance, two episodic memories of the same size may have a different number of redundant cases. Additionally, we may consider the use of case-base maintenance strategies (e.g.:[Leake, 1998; Smyth and Keane, 1995]) to improve the efficiency of the process.

From the results of the experiment, there is evidence indicating that an author-agent creates more surprising products when it is able to generate larger expectation sets. This may be explained by the fact that by generating a larger expectation set the author-agent benefits not only from the diversity of options to take into account for a specific part of a building but it also may preview the expectation set of the jury-agents with larger episodic memories and therefore with larger expectation sets. Thus, in order to build more surprising products, it may select options that are probably less expected for that jury-agents. The results do not clearly correlate the size of the expectation set of the jury-agents with the surprising-value that they assign to products, although it would be expected that the larger the expectation set, the lower the surprise felt. Actually, it would be expected that agents that expect more things for a specific part of a building would feel surprise less times. However, when the expectation set increases, the probability value is distributed by more elements (expectations) and therefore, when one of them is contradicted the surprise function that we use [Macedo and Cardoso, 2001b] computes a high surprise value. Therefore, this lack of correlation may be more related with the surprise function itself than

with the size of the expectation set. Actually, recent experiments have shown that this surprise function does not model human surprise in certain situations, namely when it computes a non-null intensity of surprise when the most expected event really happens. We are currently working on this subject.

The model for generating expectations presented in this paper makes use of the Bayes formula to compute degrees of belief (probabilities) for the expected events. After new knowledge is acquired, the probabilities of events could change reflecting the statistical data. A possible alternative approach is using Bayesian networks for this probabilistic inference (e.g., [Shafer and Pearl, 1990]). However, the techniques for learning Bayesian networks from data are still new and evolving [Heckerman and Geiger, 1995; Krause, 1998; Lam and Bacchus, 1994]. Anyway, it is certainly a promising technique worth of consideration in our work in the near future.

The theory of Boden [1992, 1995] seems to be related with our model of creativity as follows. We could establish a correspondence between creative products that, according to Boden, are to some extent impossible or "could not have happened", and those that fall in the category of practically non-deducible products. Actually, to some extent, non-deducible products could not have happened before, otherwise they would be deducible. However, we don't think that these are the only creative products, although they correspond to what is consensually accepted as the highest forms of creativity: deducible products may also be creative, although most of the times in a lower degree. The main difference is that non-deducible products are inconsistent with expectations suggested after they are perceived, i.e., they contradict normalcy, while deducible products conflict with passive or active expectations made prior to the fact. However, this conflict may be as strong as the contradiction to normalcy. Thus, according to these ideas, in our model an agent would be more creative if it is able to generate products that are beyond the expectations of the other agents, i.e., beyond the deducible knowledge of the other agents. A potential effective approach taken by an author-agent might be the construction of knowledge models of the other agents, i.e., the prediction of the deducible and non-deducible knowledge of the other agents, and subsequently the attempt to construct products that are within the non-deducible knowledge of those agents. This issue is to some extent related with Boden's concept of historical-creativity. Actually, this procedure guarantees that the created products are creative for the other agents.

5 Conclusions

A model for generating expectations seems to be crucial not only in reasoning in general but especially in creative reasoning. As shown, it is hard to imagine an agent being able to evaluate the surprise value of a product without the capability to generate expectations. The same may be applied to an agent that produces creative products. Therefore, besides the importance of knowledge (or memory) and surprise in the creativity phenomenon, there seems to be a bridge between these subjects whose relevance should not be underestimated.

Acknowledgments

We would like to thank Andrew Ortony and Rainer Reisenzein for their helpful comments. This work was partially supported by the Calouste Gulbenkian Foundation. The PhD of Luís Macedo is financially supported by PRODEP III.

References

[Boden, 1992] M. Boden. *The Creative Mind: Myths and Mechanisms*. Basic Books, New York, 1992.

[Boden, 1995] M. Boden. Creativity and unpredictability. *SEHR, 4*(2), 1995.

[Gärdenfors, 1992] P. Gärdenfors. Belief revision: An introduction. In P. Gärdenfors (Ed.) *Belief Revision*, pages 1-20. Cambridge University Press, Cambridge, UK, 1992.

[Gärdenfors, 1994] P. Gärdenfors. The role of expectations in reasoning. In M. Masuch and L. Polos (Eds.), *Knowledge Representation and Reasoning Under Uncertainty*, pages 1-16. Springer-Verlag, Berlin, 1994.

[Goel, 1992] A. Goel. Representation of Design Functions in Experience-Based Design. In D. Brown, M. Walderon and H. Yosnikawa (Eds.), *Intelligent Computer Aided Design*. Elsevier Science, 1992.

[Heckerman and Geiger, 1995] D. Heckerman and D. Geiger. Learning Bayesian Networks. Technical Report MSR-TR-95-02, Redmond, Microsoft Research - Advanced Technology Division, 1995.

[Jackson and Messick, 1967] P. Jackson and S. Messick. The person, the product, and the response: Conceptual problems in the assessment of creativity. In J. Kagan (Ed.) *Creativity and learning*, pages 1-19. Houghton Mifflin, Boston, 1967.

[Kahneman and Miller, 1986] D. Kahneman and D. Miller. Norm theory: comparing reality to its alternatives. *Psychological Review, 93*: 136-153, 1986.

[Kline, 1999] C. Kline. *Observation-based Expectation Generation and Response for Behavior-based Artificial Creatures*. MSc Thesis, Massachusetts Institute of Technology, Cambridge, USA, 1999.

[Koestler, 1964] A. Koestler. *The act of creation*. Hutchinson, London, 1964.

[Krause, 1998] P. J. Krause. Learning Probabilistic Networks. *Knowledge Engineering Review, 13*: 321-351, 1998.

[Lam and Bacchus, 1994] W. Lam and F. Bacchus. Learning Bayesian Belief Networks: an approach based on the MDL principle. *Computational Intelligence, 10*(4): 269-293, 1994.

[Leake, 1998] D. Leake, Wilson, D. Categorizing case-base maintenance: dimensions and directions. *Proceedings of the 4th European Workshop on Case-Based Reasoning*. Springer-Verlag, Berlin, 1998.

[Lubart, 1994] T. Lubart. Creativity. In R. Sternberg (Ed.) *Thinking and Problem Solving*. Academic Press, San Diego, CA, 1994.

[Macedo, 1998] L. Macedo. A Model for Creative Problem Solving based on Divergent Production of Solutions. *Proceedings of the 13th European Conference on Artificial Intelligence*. John Wiley & Sons, New York, 1998.

[Macedo and Cardoso, 2001a] L. Macedo and A. Cardoso. The Influence of the Size of the Episodic Memory on the Surprise-value of the Creative Agent's Products. *Proceedings of the Workshop on Creative Systems: Aproaches to Creativity in AI and Cognitive Science, ICCBR2001*. Naval Research Laboratory, Washington, DC, 2001.

[Macedo and Cardoso, 2001b] L. Macedo and A. Cardoso. Modelling Forms of Surprise in an Artificial Agent. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pages 588-593. Erlbaum, Mahwah, NJ, 2001.

[Macedo and Cardoso, 2001c] L. Macedo and A. Cardoso. SC-EUNE - Surprise/Curiosity-based Exploration of Uncertain and Unknown Environments. *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing*, pages 73-81. University of York, York, UK, 2001.

[Macedo and Cardoso, 2002] L. Macedo and A. Cardoso. Assessing creativity: the importance of unexpected novelty. *Proceedings of the Workshop on Creative Systems: Aproaches to Creativity in AI and Cognitive Science, ECAI2002*, pages 31-37. University Claude Bernard - Lyon, Lyon, 2002.

[MacKinnon, 1962] D. MacKinnon. The nature of nurture of creative talent. *American Psychologist, 17*: 484-495, 1962.

[Moorman and Ram, 1994] K. Moorman and A. Ram. A Model of Creative Understanding. *Proceedings of the Twelfth National Conference on Artificial Intelligence*. AAAI Press, Menlo Park, CA, 1994.

[Ortony and Partridge, 1987] A. Ortony and D. Partridge. Surprisingness and Expectation Failure: What's the Difference? *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 106-108. Morgan Kaufmann, Los Altos, CA, 1987.

[Piaget, 1952] J. Piaget. *The Origins of Intelligence in Children*. Norton, New York, 1952.

[Reisenzein, 1999] R. Reisenzein. A theory of emotions as Metarepresentational States of Mind. *Personality and Social Psychology Reviews*, 1999.

[Reisenzein, 2000] R. Reisenzein. The subjective experience of surprise. In H. Bless and J. Forgas (Eds.), *The message within: The role of subjective experience in social cognition and behavior*. Psychology Press, Philadelphia, PA, 2000.

[Ritchie, 2001] G. Ritchie. Assessing Creativity. *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, pages 3-11. University of York, York, UK, 2001.

[Russel and Norvig, 1995] S. Russel and P. Norvig. *Artificial Intelligence - A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 1995.

[Saunders and Gero, 2001] R. Saunders and J. Gero. The digital clockwork muse: a computational model of aesthetic evolution. *Proceedings of the AISB'01 Symposium on Artificial Intelligience and Creativity in Arts and Science*, pages 12-21. University of York, York, UK, 2001.

[Shafer and Pearl, 1990] G. Shafer and J. Pearl. Readings in Uncertain Reasoning. Morgan Kaufmann, Palo Alto, CA, 1990.

[Smyth and Keane, 1995] B. Smyth and M. Keane. Remembering to forget: a competence preserving case deletion policy for CBR systems. *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 377-383. Morgan Kaufmann, San Mateo, CA, 1995.

[Smyth and McKenna, 1999] B. Smyth and E. McKenna. Building compact competent case bases. *Proceedings of the Third International Conference on Case-Based Reasoning*, pages 329-342. Springer Verlag, Berlin, 1999.